

SAM (REX/Ops) Operations Meeting 2008.09.15

Adam (typist),

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

New DB server release v8.5.3 . Supports single character wild card. It is the _ . Also has changes to make rollback more robust.



sam_db_srv_pkg_v8_5_3 (10/Sep/08 | Release Notes)



SAMDEV-22 **COMPLETED** Implement single character wildcards?



SAMDEV-23 **COMPLETED** Improve db server handling of rollback failures

Eric: What about the optimized queries? CDF is running 8.4.5 - that has the most important caching optimizations. It does not have the metadata optimizations - but Robert thinks that these are less important.

No more eminent releases.

DØ:

No SAM/offline shifters during Sept. 12-14.

Sept. 9 morning (from Mike) : DZSAM-250: "CORBA errors wiped out the farm overnight";

afternoon: Andrew upgraded stager version. It should not be entering the locked state anymore.

This was the Purdue srm stager blocking problem again. We have a new stager to prevent this problem.

Sept. 10 : samfwd02 rebooted shortly before 4 pm (DZSAM-251).

Memory fried (smell of burning) on Friday. Decision 1 removed half the memory and restarted the machine (without telling anyone). This could be an important safety issue.

Sept. 12: All jobs matched to samgfw02 got screwed --> samgfw02 got burning smell

--> d0srv015 was put back into production as SAMGrid-OS-02 (DZSAM-255/257)

See above. For a short time we had two samgfw02 forwarding nodes. Steve was trying to bring up d0srv015 as SAMG-OSG-2. Real node has been running ok since

restart. Should be ok with half memory (was 4 GB, now 2 GB).

Unresolved issues:

DZSAM-54: problem on Lancaster Forwarding Node --- still being worked on (DZSAM-256 is a "spin-off").

Problem now is that Imperial College SRM won't deliver files to SAMGrid jobs - Connection Reset error. Problems seems to be SRM at Imperial College. Kostas (IC sysadmin) is away. Return date unknown.

DZSAM-236 : "single character wildcard" --- Robert has implemented a solution _ but the user probably hasn't verified.

Addressed and released. See above.

DZSAM-252: User Smaïb is asking help for metadata --- can somebody help the user to store his/her type of files ?

Not high enough priority yet. Need to give him a JIRA account.

DZSAM-254: scripts added to samgfwd01/3 --- Should be "closed" or changed into some sorts of instructions somewhere ??

Basic init script is there (for root), need to add the samgrid and sam user scripts.

DZSAM-255/257: Problems with samgfwd02 --> d0srv015 as replacement (for now).

See above.

Also problem with old jobs trying to get files from dead samgfwd02 -- fcp was retrying forever. Version in samgrid does have a total timeout (Parag says that there is a 3600 second timeout, but we never saw that work - timeout is specified in jim_job_managers or sam_fcp_config). The newer version of fcp does have a connection level timeout.

DZSAM-256 : "File transfer problem between d0srv071 and Imperial College SRM" --- being actively worked on.

See above

DZSAM-258: "Add new stage area for OSG production at MIT" --- very new request for a (new?) SRM station.

Done - added a new SRM machine type.

CDF:

Looked like some LTO4 drives went offline - but it was really a monitoring problem. e.g. if you click on the link and no drives are listed, then monitoring timeout (these are on the enstore web page). Must be careful about this.

Thursday - Alex tried again to reconfig the dCache pools - now have all read pools (no more stager and read-only pools).

fcdfdata147 was dead but is now back (pool node - motherboard and memory were replaced).

Thursday night/Friday morning - Network outage - (Art:) the primary DNS server stopped providing services.

Friday night - ran out of disk space on cdfsamweb (old one). So web pages were incomplete - online thought they lost their db server. Angela cleaned up the disk.

MINOS:

Catalogued several problems during last Friday's DNS outage

Helpdesk down - could not submit tickets, or do Expert login

MRTG monitoring stopped

DocDB down

FNDCA/STKEN stopped delivering data

Control Room Logbook failed - due to our mysql host resolv.conf error

Not clear why the other two backup DNS servers didn't pick up the service?

Why didn't enstore use the #2, #3 DNS servers?

Enstore -

DCache -

SAM -

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

A disk purchase in MISER was shown as being signed off, but that was not the case.

Be careful in MISER -- you can get the lab req number (click on title under icon - will be made better). This lead to a couple of month delay.

DB:

D0 :

Alert(s)

- Tablespace [OPERATIONAL_01] is [80 percent] full

Action: Added 4GB

Also added 6GB to OPERATIONAL_IDX01

CDF:

- OS patches were deployed on **cdf** offline dev and int host node.
- Scheduling downtime for cdfopr2 for OS & DB security patches for Sep 18.
Also We will set the explicit limit for sessions to 900 versus 665. This will help to fix the intermittent ORA-018 (excessive number of sessions) error.

Downtime 8:30 - noon.

Alert(S):

- Information for: cdfopr2

Owner Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD ANALYSIS_PROCESSES	TABLE	196	200	4

Action: Increased the max extents to 500

MINOS:

Alert(s):

Enstore/dCache:

From CD Ops report.

* DSC section operational issues:

Enstore Operations report

- Release Management:

+ Scheduled Downtimes:

CDFen downtime: Thursday 9/18 at 8:30 AM Upgrade the Movers and Library Manager code. Upgrade fi rmware on raids.

- Problem Management:

+ Unscheduled Downtimes:

None.

- Capacity Management:

+ Agilysys has been awarded the the SL8500 purchase of 2 SL8500's.
Detail for delivery and install have not been finalized. The fire suppression system POs have NOT yet been cut (\$75 on Run II and \$75K CMS tasks)

+ Transferring D0 LTO4 tapes from SL8500#2 (GCC) to SL8500#3 (FCC).
These are non-raw data tapes as part of the plan to store Run II RAW at GCC and Run II non-raw at FCC

- Incident Management:

Both stkendca6a and 7a both have bad power supplies and are being replaced.
Reinstalled several movers for disk issues.

Data Integrity:

VSN PRQ585. Still waiting for Sun to respond.

Historical report (a work in progress):

7/14 7/21 7/28 8/4 8/11 8/18 8/25 9/8 9/15

17 | 21 | 48 | 5 | 25 | 4 | 36 | 19 | 25 tape mover interventions
20 | 34 | 50 | 30 | 24 | 6 | 0 | 48 | 20 tape incidents
1 | 1 | 1 | 0 | 2 | 1 | 3 | 0 | 0 library hw incidents
0 | 0 | 1 | 1 | 0 | 1 | 0 | 2 | 0 tape drive hw incidents
0 | 0 | 0 | 5 | 1 | 0 | 0 | 0 | 7 enstore mover hw incidents
0 | 0 | 2 | 1 | 0 | 1 | 1 | 0 | 0 enstore server hw incidents
2 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 3 | 4 dcache pool hw incidents
0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 dcache server hw incidents
0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 pnfs hw incidents
3 | 3 | 0 | 3 | 0 | 1 | 0 | 0 | 0 | 0 dcache sw incidents
0 | 1 | 0 | 2 | 0 | 1 | 1 | 1 | 0 | 0 | 1 enstore sw incidents
0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 enstore conf incidents
1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 3 dCache conf incidents

7 | 4 | 0 | 0 | 2 | 0 | 12 | 0 | 0 | 0 | 0 unsched enstore outage hrs
0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 unsched dCache outage hrs

For plots of above see: http://www-ccf.fnal.gov/IA_metrics/

- Migration work flow status:

LTO1 Migration: 0 / ~1200
LTO2 Migration: 2714 / ~8500
9940A Migration: 1728 / ~4500
9940B Migration: 0 / ~24000
LTO1 Migration: 0 / ~1200

http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_daily_3480.jpg

http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_accumulated_3480.jpg

http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_daily_9940.jpg

http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_accumulated_9940.jpg

- SSA group is interviewing for a full time position. Interviews and phone screens in progress.

Public dCache:

* Pool node stkendca11a has been restarted at 9:15 AM on Wednesday, one pool hasn't started up. The PID file has been removed, pool restarted.

* Pool node stkendca13a was power cycled by accident on Thursday morning. Two pools haven't started up because of the same reason - PID files did exist and system believed the pools were running. Pool stop/start cycle solved the problem.

CDF dCache:

* On Wednesday morning pool manager configuration was changed per experiment request to convert readonly pools and stage pools to read pools. It resulted in interruption to file restores. The configurations were rolled back for investigation. The cause was a missing link to real-encp script on former read-only pool. The link was fixed, the new configuration reloaded, and is now operational.

* Node fcdldata147 returned into the system after having some memory replaced. Once again two pools didn't start up because the same problem with PID files. Alex K. took care of it, the pools have been restarted.

* Four read pools and one raw pool on nodes fcdldata318 and fcdldata325 are offline. Investigating.

Special: Power outage on 9/18 at GCC

- Thursday, September 18th -- GCC Power Outage
- 480V repair, testing, maintenance - **no generators**
- CRA 1-3 hours, CRB 2-8 hours (CRA = computer room A)
- will bring cooling switch boards back first
- expect GCC networking & tape libraries to remain up on UPS
- GCC/CRC and LCC 480 maintenance tbd (Sept. 17, 18, 19)
- try to schedule LCC 13kV maintenance with LCC 480V maintenance
- date still tentative based on availability of FESS resources